

(19)



Europäisches Patentamt
European Patent Office
Office européen des brevets



(11)

EP 0 901 307 A2

(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:
10.03.1999 Bulletin 1999/10

(51) Int Cl.⁶: H04Q 11/04, H04L 12/56

(21) Application number: 98307144.0

(22) Date of filing: 04.09.1998

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE
Designated Extension States:
AL LT LV MK RO SI

(30) Priority: 05.09.1997 US 924657

(71) Applicant: NEWBRIDGE NETWORKS
CORPORATION
Kanata, Ontario K2K 2E6 (CA)

(72) Inventors:
• Carr, David Walter
Nepean, Ontario, K2J 3Z5 (CA)
• Lee, Denny L.S.
Kanata, Ontario, K2K 2P8 (CA)
• Davis, Tom
Ottawa, Ontario, K1Y 4B9 (CA)

(74) Representative: Harding, Richard Patrick et al
Marks & Clerk,
4220 Nash Court,
Oxford Business Park South
Oxford OX4 2RU (GB)

(54) Virtual path shaping

(57) Apparatus and method for shaping ATM cell traffic emitted onto a virtual path connection in an ATM network are described. Component virtual channel connections are arbitrated at an aggregation point utilizing

an arbitration technique. The technique provides both virtual path shaping and controllability of underlying virtual channel connections with an improved fairness performance amongst all the aggregating virtual channel connections.

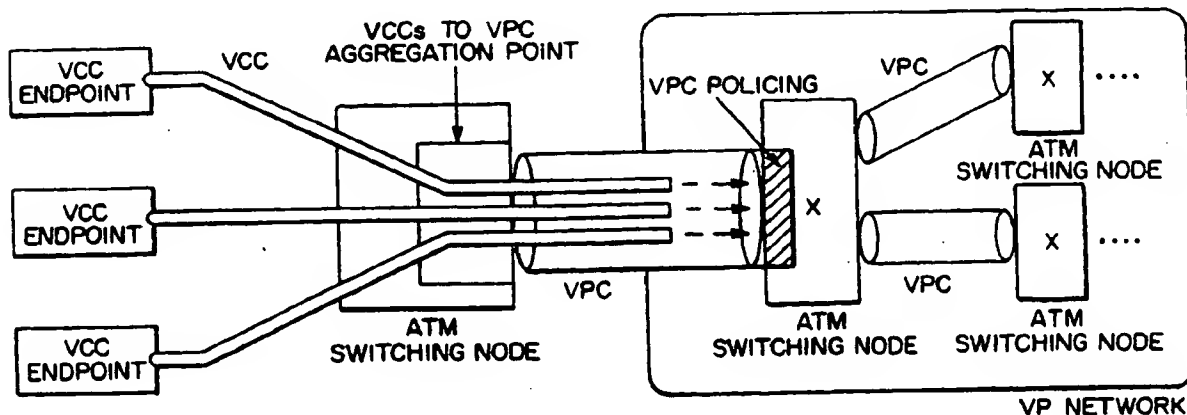


FIG. 3

Description

Field of the Invention

[0001] This invention relates to data communications networks utilizing ATM technology and more particularly to apparatus and methods for scheduling/shaping multi-service category cell traffic onto virtual path connections while providing fair share arbitration between aggregating virtual channel connections.

Background of the Invention

[0002] Asynchronous transfer mode (ATM) technology is rapidly being recognized as the technology of choice for broadband data communications. ATM cells having a fixed number of bytes can efficiently accommodate multimedia traffic such as video, voice and data.

[0003] ATM cell traffic between nodes in an ATM network is carried on what are known as virtual connections (VC). Traffic from end user to end user through a plurality of intermediate nodes is carried on virtual channel connections (VCC). A VCC on a particular link is uniquely identified by the combination of virtual path identifier (VPI) and virtual channel identifier (VCI) fields in the ATM cell headers. Since traffic from several sources may have a common destination it is known to combine virtual channel connections onto a virtual path connection (VPC) in order to simplify the processing of traffic through the network. A VPC on a particular link is uniquely identified by the VPI field in the ATM cell header. When virtual channel connections are aggregated onto a virtual path connection, the individual virtual channel connections are essentially transparent to the intermediate ATM nodes and this allows a simpler networking model and higher manageability.

[0004] A virtual path connection is defined as a logical grouping of multiple virtual channel connections into one virtual connection. An implicit logical aggregation of VCCs into a VPC is said to have occurred when the outgoing VCCs all acquire the same common VPI value. Typically, such aggregation occurs at the egress port of an ATM switch where multiple VCCs from multiple ingress ports converge. These aggregating VCCs might have different VPI and VCI values as determined at call setup or network planning time. The egress port of the ATM switch is typically equipped with a cell header translation function. Such a function is capable of translating the incoming VPI and VCI values into a new set of VPI and VCI values. It is this function that accomplishes the aforementioned implicit logical aggregation. When the various VPI and VCI values on the various aggregating VCCs all translate into a common VPI value, a virtual path connection is formed. Note that the VCI value is carried transparently end-to-end within a VPC. It is, therefore, important for the translation function to translate the incoming VCI values into a mutually exclusive set of VCI values. Such mutual exclusivity al-

lows each underlying VCC to be uniquely identified at the terminating end of the VPC. At that point, the VPC is segmented by simply switching and processing each underlying VCC as an individual connection. The combination of an egress aggregation function together with the ingress segregation function on a bi-directional port is referred to as the virtual path termination point (VPTP).

[0005] Although the aggregation of virtual channel connections onto a common path simplifies traffic management, concerns have been expressed that individual virtual channel connections may not all be treated fairly when aggregated onto a virtual path connection. Demirtjis and Petr (IEEE IC3N, 1995 "How to Encourage Cheating with Virtual Path Policing in ATM Networks") have found that virtual channel connections which are in compliance with predefined transmission parameters may be adversely affected by non-compliant traffic when aggregated onto a virtual path.

[0006] US Patent 5,629,928 which issued May 13, 1997 to Calvignac et al discloses one technique for controlling traffic flow in a virtual path communication system wherein queues allocated as needed from a pool of free queues are used to store all incoming data cells having the same virtual channel identifier.

Summary of the Invention

[0007] The present invention seeks to provide shaping of virtual path traffic through an ATM network utilizing an arbitration technique.

[0008] Therefore, in accordance with a first aspect of the present invention there is provided a device for scheduling virtual channel connections carrying cell traffic of different categories of service onto a virtual path connection comprising: aggregating means for aggregating virtual channel connections of similar categories of service into a single cell stream.

[0009] In accordance with a second aspect of the present invention there is provided, in an ATM data communications network utilizing virtual channel connections wherein the virtual channel connections are aggregated at an aggregation point onto virtual path connections, an apparatus for shaping ATM cell traffic comprised of various categories of service onto the virtual path connection, the apparatus comprising: queuing means for queuing cells from each of the virtual channel connections; arbitration means to arbitrate between the virtual channel connection traffic; and a virtual path connection shaper to determine egress emission time of cells from the arbitration means onto the virtual path connection.

[0010] In a preferred embodiment of this aspect of the invention virtual channel connections carrying the various categories of service, as defined by the ATM Forum are queued separately. Like-service-category virtual channel connection queues are grouped into pe groups. An arbitration scheme arbitrates within

per-VC group and another arbitration scheme arbitrates between per-VC groups with like categories of service

[0011] In accordance with a further aspect of the present invention there is provided in an ATM data communications network having virtual channel connections channel aggregated at an aggregation point onto virtual path connections, a method for fairly shaping ATM cell traffic comprised of various categories of service and quality of service parameters onto the virtual path connection, the method comprising the steps of: queuing the virtual channel connection cell traffic in separate queuing buffers; arbitrating between the separate queuing buffers utilizing an arbitration means; and shaping the arbitrated cell traffic with a virtual path connection shaper in order to fairly emit the cell traffic onto the virtual path connection.

[0012] In a preferred embodiment of this aspect of the invention the shaper determines the cell egress times in accordance with certain traffic descriptors.

Brief Description of the Drawings

[0013] The invention will now be described in greater detail with reference to the attached drawings wherein:

Figure 1 is one prior art method of forming a virtual path connection at an aggregation point;

Figure 2 represents a second prior art method of forming a virtual path connection at an aggregation point;

Figure 3 is a block diagram of an ATM network with VPC aggregation;

Figure 4 illustrates the virtual path shaping technique of the present invention.

Figure 5 is a functional block diagram of the vertical path shaping process;

Figure 6 is a flow diagram of the process depicted in Figure 5.

Figure 7 illustrates one real time virtual path aggregation scheme; and

Figure 8 illustrates an alternate aggregation scheme.

Detailed Description of the Invention

[0014] Known virtual path connection (VPC) aggregation is typically performed at the egress direction of an ATM switch. Such VPC aggregation points do not usually have VP level shaping capability and, therefore, the aggregated cell stream cannot be policed in downstream nodes. Other VPC aggregation points may have

per-VC shaping capability, therefore the VP must first be aggregated before it is shaped. This method does not allow controllability of the underlying VCCs at the VP aggregation point. These controllable elements include such items as per-VC buffer management, rate-control and fairness amongst aggregating VCCs.

[0015] Figure 1 illustrates a logical aggregation at a VPC aggregation point. As shown, the component VCCs are multiplexed onto the virtual path connection without any virtual path shaping. This results in an uncontrolled output cell stream which is unlikely to provide fair sharing of the virtual channel connection cells. Without shaping, the aggregated VPC traffic is likely to include irregular flow.

[0016] Figure 2 illustrates a VPC aggregation point with shaping after aggregation. As in Figure 1 the component VCCs are aggregated at the input to the VPC. A single aggregate queue stores the incoming component VCCs prior to the per-VC shaper which shapes traffic on the VPC. There is, however, no attempt to provide quality-of-service differentiating arbitration between the component VCC traffic.

[0017] The technique contemplated by the present invention provides VP aggregation shaping while maintaining per-VCC isolation amongst the aggregating VCCs. Such isolation makes it possible to maintain per-VC congestion control and quality of service differentiation on each individual component VCC. For the purpose of this description the terms "component" and "aggregating" are used interchangeably. They both denote the underlying VCCs which constitute a VPC. A network block diagram is shown in Figure 3.

[0018] The ATM Forum has defined five service categories relating to ATM cell traffic. The five service categories are: constant bit rate (CBR), real time variable bit rate (rt-VBR), non-real time variable bit rate (nrt-VBR), available bit rate (ABR) and unspecified bit rate (UBR). These categories can be divided into two main groups: real time (CBR and rt-vbr) and non-real time (nrt-VBR, ABR and UBR). For the real time service categories cell transfer delay (CTD) and cell delay variation (CDV) are both important quality of service (QoS) parameters. In light of these considerations the present invention relates to a hierarchical approach in which individual VCCs are logically grouped together with common service requirements. The technique for virtual path shaping according to the present invention is illustrated in Figure 4.

[0019] As shown in Figure 4, each component VCC is individually queued in its respective per-VCC queue. As such any buffer management or congestion control technique can be applied to the isolated VCCs. This feature inherently allows better fairness as each connection's behaviour is isolated from one another. In view of the aforementioned definition of multiple service categories wherein each requires a different quality of service, the present invention addresses the relevant issue by introducing multiple arbitration hierarchies. The low-

est level depicted in Figure 4 is the service group level. All component VCCs within the same service group effectively share the same level of service as one unique arbitration technique is applied to all members of a particular service group. The arbitration can be, for example, a simple round-robin technique, an exhaustive round-robin technique, a weighted round-robin technique, a weighted fair queuing technique or any other scheduling technique. The arbitration choice is dependent on the performance and implementation complexity desired. For the best fairness performance a per-VC weighted fair queuing is used. This technique, however, incurs the highest level of implementation complexity.

[0020] After the VCCs have been arbitrated as between service groups the next hierarchical level is to arbitrate amongst all the service groups. As a potential application of this technique it is possible to map the five service categories defined by the ATM Forum into five separate service groups. Otherwise, it is possible to map the real time service categories into one group and the non-real time service categories into the other. The usage and partition of the service groups are considered to be implementation decisions. Depending on the desired level of complexity versus performance, it is possible to choose from the various arbitration techniques discussed previously. For instance, if the service groups are mapped into real time and non-real time service categories, the inter-group arbitration can be as simple as exhaustive servicing with the real time group having absolute priority over the non-real time group. This gives the real time group's connections better delay performance.

[0021] The two levels of arbitration determines the next cell which is to be transmitted and the shaper determines in the time domain when that cell can be emitted for the virtual path connection given certain traffic descriptors. These descriptors might include peak cell rate (PCR), sustained cell rate (SCR), and maximum burst size (MBS). Upon the determination that a cell can be emitted the inter-group's arbiter determines which service group is allowed to send and the chosen service group arbiter then further determines which component VCC should be served. Once the VCC is determined the cell at the head of that per-VCC queue is extracted and transmitted.

[0022] With this VP aggregation technique it is possible to achieve both VP shaping and controllability of underlying VCCs with an improved fairness performance amongst all the aggregating VCCs.

[0023] Figures 5 and 6 provide, respectively, a function diagram and a flow diagram of the shaping process.

[0024] Figure 5 sets out the following functional steps:

- 1) cell arrival event;
- 2) arrival notice;
- 3) discard decision based on current congestion condition;
- 4) queue manager determines the physical enqueue-

ing location;

- 5) VPC shaper scheduling cell emission in time domain;
- 6) VCC arbitration block signals the chosen component VCC given a certain arbitration scheme;
- 7) queue manager dequeues the cell from the chosen component VCC; and
- 8) cell transmission

[0025] Figure 6 shows the steps in flow diagram format.

[0026] The system according to the present invention supports aggregation for all of the aforementioned service categories. An end-to-end VPC connection originates at an egress port (near-end node) and ends at an ingress port (far-end node). These ports or points are denoted VP Termination Points (VPTPs). Accordingly, a VPC aggregates at egress and segregates at ingress. The system support for VPCs pertains to its ability to shape and arbitrate amongst the aggregating connections. A description of an architectural model for arbitrating real-time service categories follows.

[0027] For real-time VPC aggregation in egress, the system supports aggregation of all five of the previously identified service categories. Figures 7 and 8 graphically illustrates two options for the handling of priorities within the aggregating VPTP.

[0028] The first option is shown in Figure 7. As shown Universal Card Slots (UCSs) 30 are connected via VCCs 32 through the switching fabric 34 to the arbitration/shaper device 36 of this embodiment. Device 36 includes per-VC queues 38 (CBR/rt-VBR), 40 (nrt-VBR), 42 (ABR) and 44 (UBR). Around-robin (RR) arbitration scheme 46 is implemented within each VC group. Hierarchical arbitration as between service categories is implemented by Exhaustive Round Robin scheme 48 and Weighted Round Robin scheme 50 as will be discussed later. Temporal shaper 52 determines egress times of cells from respective queues to VPC 54.

[0029] In the egress direction, aggregating VCs are per-VC queued. CBR and rt-VBR categories of service are serviced exhaustively over all other service categories. These CBR and rt-VBR connections are serviced in a round-robin (RR) fashion amongst themselves. For non-real time (nrt) VBR, ABR and UBR connections, they are serviced in a weighted-round-robin (WRR) fashion. As in the real time case like-service-categories connections are serviced in a simple round-robin fashion.

[0030] The reason for separating nrt-VBR, ABR and UBR with WRR is as follows: nrt-VBR is separated out so that sustainable cell rate (SCR0) and peak cell rate (PCR0 and PCR1) type shaping would not take up all the bandwidth if users choose to send cell loss priority (CLP1) cells. ABR and UBR are separated to have their own weight so that a simple explicit rate (ER) algorithm can be used and still achieve available bandwidth partitioning between ABR and UBR in both port and VPTP

levels.

[0031] Given that usage of CLP1 traffic in VBR service category can be ignored taking into account usage of the ER algorithm, a simpler scheme for VPC aggregation can be implemented as shown in Figure 8. As shown in Figure 8 there are three groups of per-VC queues, CBR and rt-VBR 60, nrt-VBR 62, and ABR/UBR 64. These three groups are serviced in an exhaustive-round-robin fashion (66). Within each group, they are serviced with a simple round-robin (RR) scheme. The nrt-VBR traffic is isolated such that UBR traffic cannot overwhelm nrt-VBR connections.

[0032] The outgoing shaping priority and characteristics are associated with the outgoing VPC's service category. With the per-VC queuing nature maintained, the buffer management strategy operates similarly with respect to other connection types (i.e. simple VCC, VPC switching connections) Discard algorithms such as early packet discard (EPD) work on a per-VC basis for these aggregating VCCs.

[0033] For non-real-time VPCs, all underlying traffic must be non-real-time. Therefore, there is no exhaustive-round-robin arbitration between the real-time group and the non-real-time group. The WRR scheme still applies to the aggregating non-real time VCC in the option of Figure 7. For the option of Figure 8 the CBR/rt-VBR group is simply removed from the inter-group contention and the ER scheme is replaced with WRR.

[0034] While particular embodiments of the invention have been disclosed and illustrated, it will be apparent to one skilled in the art that variations and alternatives can be implemented without departing from the basic concept of the invention. It is to be understood that such variations and alternatives, to the extent possible, fall within the scope of the invention as defined by the appended claims.

Claims

1. A device for scheduling virtual channel connections carrying cell traffic of different categories of service onto a virtual path connection comprising:

aggregating means for aggregating virtual channel connections of similar categories of service into a single cell stream.
2. A device as defined in claim 1 wherein virtual channel connections carrying different categories of service are queued in separate queues.
3. A device as defined in claim 2 wherein said aggregating means includes means to form groups of said separate queues
4. A device as defined in claim 3 wherein said aggregating means includes arbitration means.
5. A device as defined in claim 4 wherein said aggregation means includes means to form per-VCC like-service-category queues and said arbitration means arbitrates priority as between said like-category-queues.
6. A device as defined in claim 5 wherein said arbitration means includes means to arbitrate priority as between said separate queues within said groups of queues.
7. A device as defined in claim 1 including multilevel arbitration means to arbitrate between virtual channel connections in an hierarchical manner.
8. A device as defined in claim 6 including shaper means to determine egress time of cells in said single cell stream.
9. A device as defined in claim 8 wherein said egress time is shaped in accordance with traffic descriptors specified for said virtual path connection.
10. A device for shaping virtual channel connections carrying cell traffic of different categories of service onto a virtual path connection comprising: queuing means for queuing cells of like-service-categories; arbitration means for arbitrating between said queues; and shaper means for shaping said cells into a single cell stream on said virtual path connection.
11. In an ATM data communications network utilizing virtual channel connections wherein said virtual channel connections are aggregated at an aggregation point onto virtual path connections, an apparatus for shaping ATM cell traffic comprised of various categories of service onto said virtual path connection, said apparatus comprising:

queuing means for queuing cells from each of said virtual channel connections;

arbitration means to arbitrate between said virtual channel connection traffic; and

a virtual path connection shaper to determine egress emission time of cells from said arbitration means onto said virtual path connection.
12. An apparatus as defined in claim 11 wherein said virtual channel connections carry different categories of service cell traffic.
13. An apparatus as defined in claim 12 wherein said queuing means queues said virtual channel connections according to said category of service.
14. An apparatus as defined in claim 13 wherein said arbitration means arbitrates priority as between

said virtual channel connections.

15. An apparatus as defined in claim 14 further including second arbitration means to arbitrate priority as between categories of service 5
16. An apparatus as defined in claim 15 wherein said categories of service comprise real time and non-real time service classes 10
17. An apparatus as defined in claim 16 wherein said real time class of service include constant bit rate (CBR) and real time variable bit rate (rtVBR) traffic.
18. An apparatus as defined in claim 16 wherein said non-real time class of service includes non-real time variable bit rate (nrtVBR), available bit rate (ABR) and unspecified bit rate (UBR). 15
19. An apparatus as defined in claim 11 wherein said shaper determines egress time of cells in accordance with traffic descriptors. 20
20. An apparatus as defined in claim 19 wherein said traffic descriptors include peak cell rate (PCR), sustained cell rate (SCR) and maximum burst size (MBS). 25
21. An apparatus as defined in claim 15 wherein said second arbitration means utilizes an exhaustive round-robin arbitration scheme. 30
22. An apparatus as defined in claim 15 wherein said second arbitration means utilizes a weighted round-robin arbitration scheme. 35
23. An apparatus as defined in claim 15 wherein said second arbitration means utilizes a weighted fair queuing arbitration scheme. 40
24. In an ATM data communications network having virtual channel connections channel aggregated at an aggregation point onto virtual path connections, a method for fairly shaping ATM cell traffic comprised of various categories of service and quality of service parameters onto said virtual path connection, said method comprising the steps of: 45

queuing said virtual channel connection cell traffic in separate queuing buffers; 50
 arbitrating between said separate queuing buffers utilizing an arbitration means; and
 shaping said arbitrated cell traffic with a virtual path connection shaper in order to fairly emit said cell traffic onto said virtual path connection. 55

25. A device for arbitrating and shaping virtual channel

connections (VCC) onto a virtual path connection (VPC) wherein said virtual channel connections carry ATM cell traffic of different categories of service, said device comprising: queuing means for forming per-VCC queues containing similar category of service cells; first arbitration means within each per-VCC queue for arbitrating priority between similar category of service cells; a second arbitration means for arbitrating priority between said per-VCC queues; and shaper means for determining egress time of said cells from said second arbitration means onto said virtual path connection

26. A device as defined in claim 25 wherein a third arbitration means arbitrates between some of said per-VCC queues with the results of said third arbitration means being arbitrated by said second arbitration means together with the other per-VCC queues. 15
27. A device as defined in claim 26 wherein said first arbitration means is a round robin scheme. 20
28. A device as defined in claim 26 wherein said second arbitration means is an exhaustive round robin scheme. 25
29. A device as defined in claim 26 wherein said third arbitration means is a weighted round robin scheme. 30




FIG. 1




FIG. 2




FIG. 3

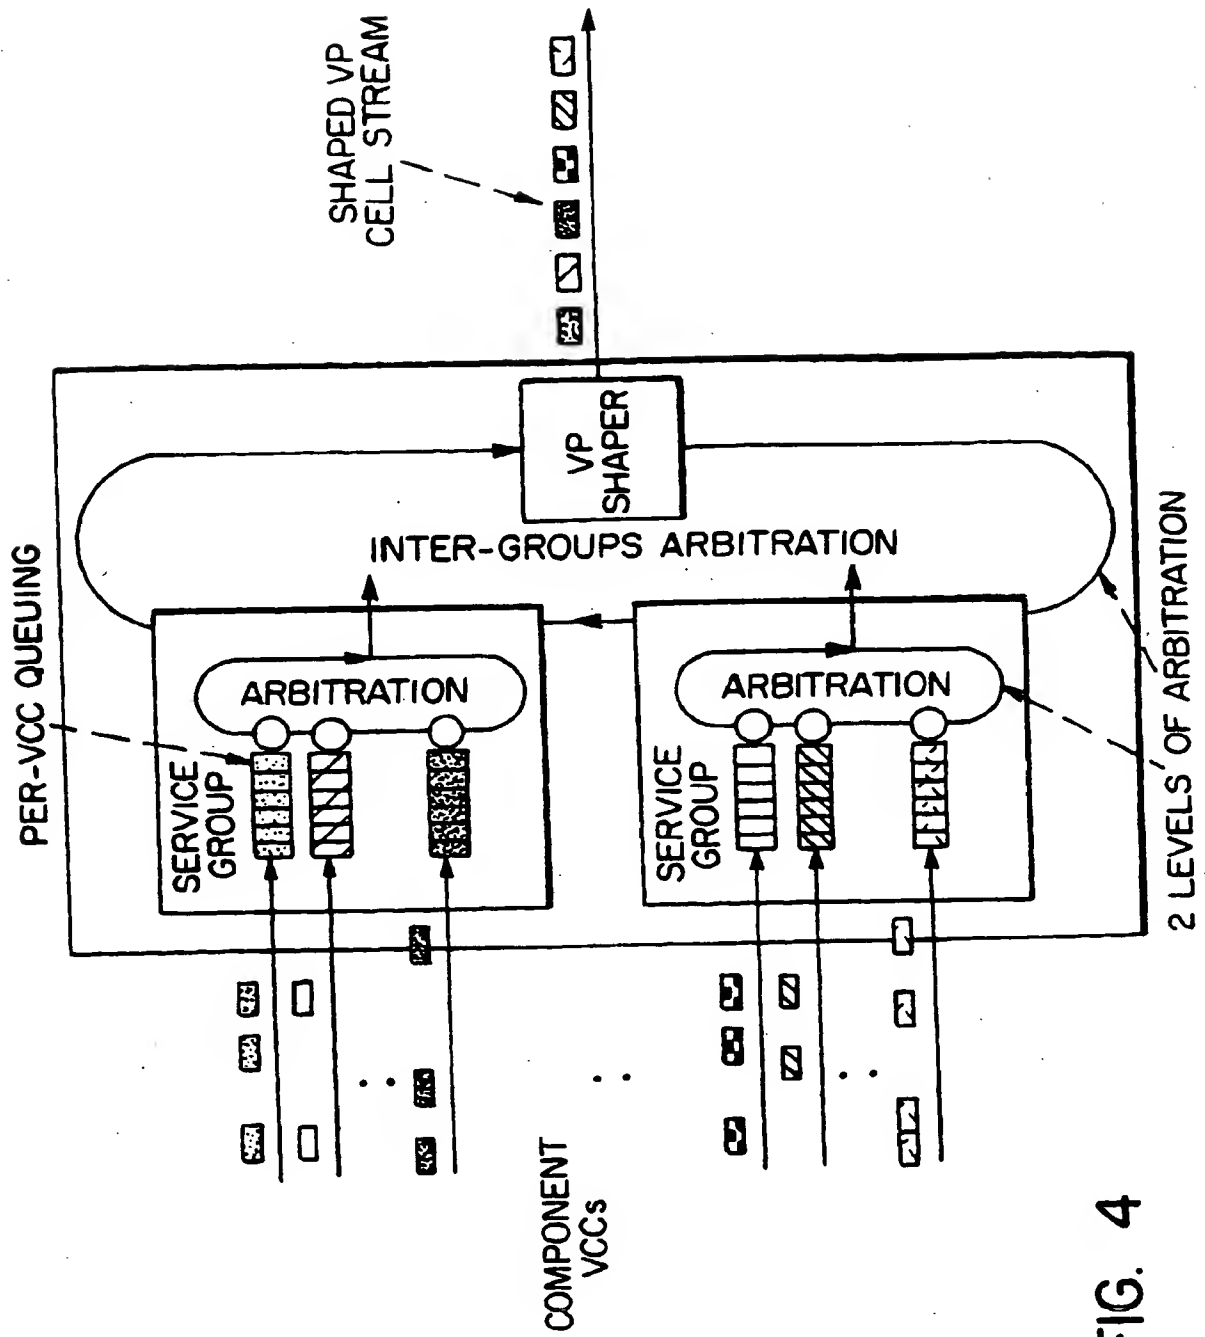


FIG. 4


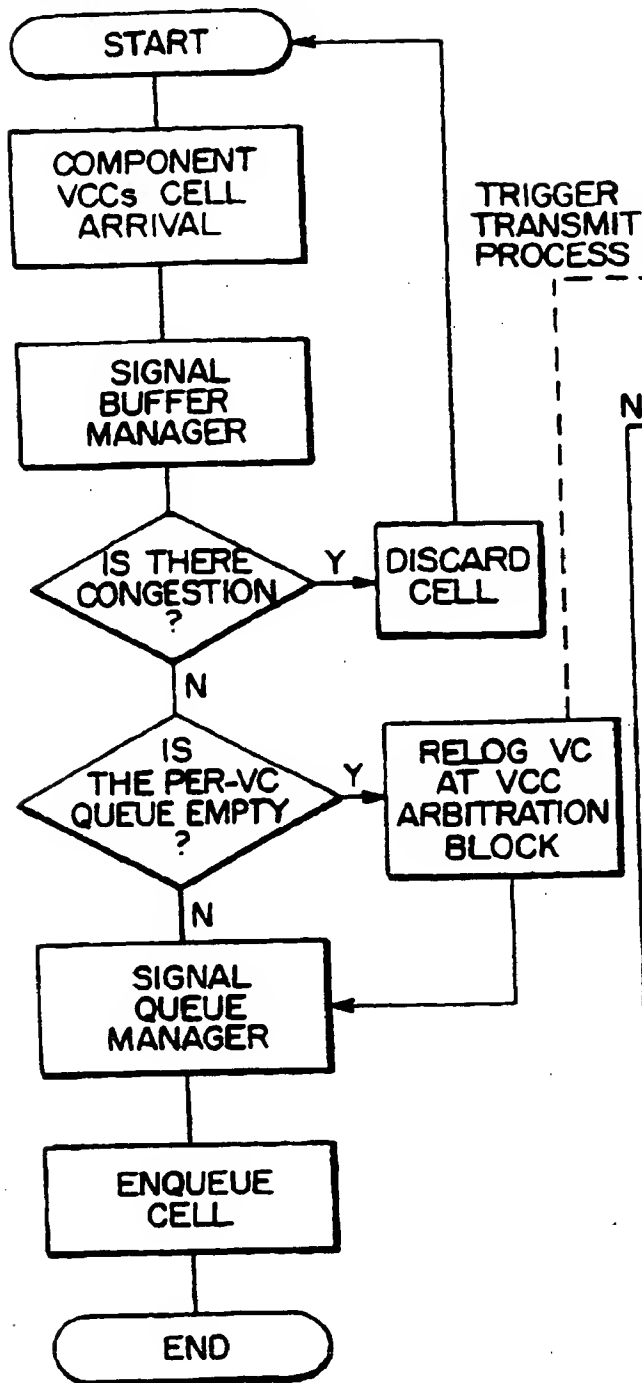


FIG. 5

RECEIVE PROCESS



TRANSMIT PROCESS

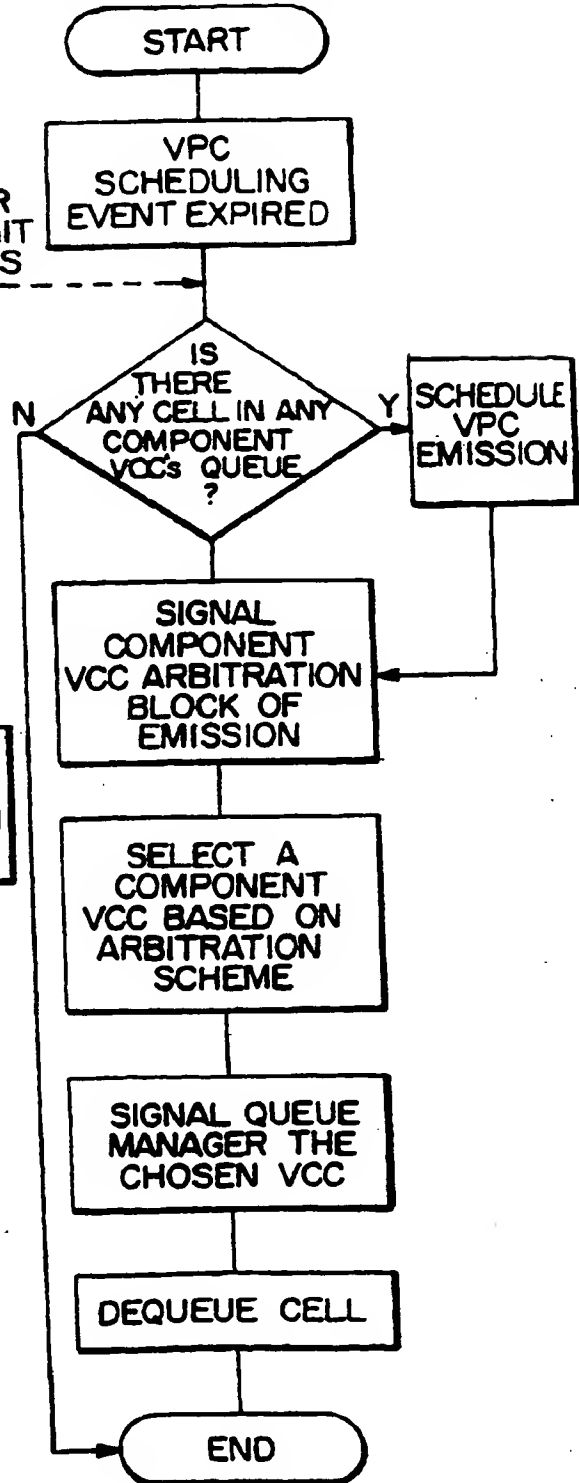


FIG. 6




FIG. 7




FIG. 8

(19)



Europäisches Patentamt
European Patent Office
Office européen des brevets



(11)

EP 0 901 307 A3

(12)

EUROPEAN PATENT APPLICATION

(88) Date of publication A3:
08.03.2000 Bulletin 2000/10

(51) Int Cl.7: H04Q 11/04, H04L 12/56

(43) Date of publication A2:
10.03.1999 Bulletin 1999/10

(21) Application number: 98307144.0

(22) Date of filing: 04.09.1998

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE
Designated Extension States:
AL LT LV MK RO SI

(72) Inventors:
• Carr, David Walter
Nepean, Ontario, K2J 3Z5 (CA)
• Lee, Denny L.S.
Kanata, Ontario, K2K 2P8 (CA)
• Davis, Tom
Ottawa, Ontario, K1Y 4B9 (CA)

(30) Priority: 05.09.1997 US 924657


(71) Applicant: **NEWBRIDGE NETWORKS
CORPORATION**
Kanata, Ontario K2K 2E6 (CA)

(74) Representative: **Harding, Richard Patrick et al**
Marks & Clerk,
4220 Nash Court,
Oxford Business Park South
Oxford OX4 2RU (GB)

(54) Virtual path shaping

(57) Apparatus and method for shaping ATM cell traffic emitted onto a virtual path connection in an ATM network are described. Component virtual channel connections are arbitrated at an aggregation point utilizing

an arbitration technique. The technique provides both virtual path shaping and controllability of underlying virtual channel connections with an improved fairness performance amongst all the aggregating virtual channel connections.

**FIG. 3**



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 98 30 7144

DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.6)
X	EP 0 596 624 A (AMERICAN TELEPHONE & TELEGRAPH) 11 May 1994 (1994-05-11) * page 2, left-hand column, line 55 - right-hand column, line 37 *	1-29	H04011/04 H04L12/56
X	HLUCHYJ M G ET AL: "QUEUING DISCIPLINES FOR INTEGRATED FAST PACKET NETWORKS" PROCEEDINGS OF THE INTERNATIONAL CONFERENCE ON COMMUNICATIONS, US, NEW YORK, IEEE, vol. -, page 990-996 XP000326820 ISBN: 0-7803-0599-X * page 994, left-hand column, line 50 - right-hand column, line 37 *	1-29	
A	ROBERTS J W: "VIRTUAL SPACING FOR FLEXIBLE TRAFFIC CONTROL" INTERNATIONAL JOURNAL OF COMMUNICATION SYSTEMS, GB, WILEY, CHICHESTER, vol. 7, no. 4, page 307-318 XP000672879 ISSN: 1074-5351 * page 310, left-hand column, line 42 - page 312, right-hand column, line 20 *	1-29	
			TECHNICAL FIELDS SEARCHED (Int.Cl.6)
			H04Q H04L
The present search report has been drawn up for all claims			

Place of search

THE HAGUE

Date of completion of the search

7 January 2000

Examiner

Veen, G

CATEGORY OF CITED DOCUMENTS

- X : particularly relevant if taken alone
- Y : particularly relevant if combined with another document of the same category
- A : technological background
- O : non-written disclosure
- P : intermediate document

- T : theory or principle underlying the invention
- E : earlier patent document, but published on, or after the filing date
- D : document cited in the application
- C : document cited for other reasons
- S : member of the same patent family, corresponding document

This Page Blank (uspto)